Probabilistic Hazard Detection for Autonomous Safe Landing

Tonislav Ivanov
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
818-354-5017
tivanov@jpl.nasa.gov

Andres Huertas
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
818-393-6695
huertas@jpl.nasa.gov

Abstract — **Future generation of landing craft will autonomously look at the surface during the terminal phase of powered descent and then, in real-time, choose and divert to a safe landing site in order to avoid hazards. Enabling technologies for such capability have been under development in recent years in the Autonomous Landing Hazard Avoidance Technology (ALHAT) project funded by NASA's Exploration Technology Development Program. ALHAT is a comprehensive system that spans the approach and landing events – from de-orbit coasting to touchdown. In this paper, we focus on ALHAT's perception task of detecting hazards in the sensed terrain and of selecting candidate safe sites for landing. This task, named Hazard Detection and Avoidance (HDA), occurs in the middle of the landing sequence. Our approach to HDA employs a probabilistic model in order to better manage the ubiquitous uncertainties associated with noisy sensor measurements and navigation. Also, we explicitly take into account the geometry of the lander and its interaction with the surface when assessing hazards. Experimental results on synthetic Lunar-like terrain show that our HDA algorithm can designate safe landing locations for a variety of terrain types and density and abundance of hazards. The complete ALHAT system is undergoing ground field-testing, and is scheduled for additional field tests on a one-hectare, lunar-like, hazard field recently constructed at NASA's Kennedy Space Center (KSC). Although the focus of ALHAT is on autonomous planetary landings, a number of terrestrial applications can also benefit from out HDA system.**

TABLE OF CONTENTS

1. INTRODUCTION

Landing of spacecraft requires autonomy. For unmanned missions, this need is essentially due to the time delay in communication between the spacecraft and operators on Earth. For manned missions, to the Moon for example, autonomy can assist astronauts when landing in dark or hazardous regions with minimal human supervision required. However, autonomous landing remains very challenging; it is probably the most critical engineering aspect of any mission. It is not for no reason that the final moments of Mars landings have been correctly dubbed "minutes of terror." Thus, the development of safe landing capabilities, such as on-board HDA, is strongly desired. In fact, robotic landings with automated HDA are among the goals of NASA, other nations' space agencies, and even private companies. Current efforts are underway in NASA's ALHAT and Lander Vision System (LVS) projects to design, build, and validate HDA systems up to Technology Readiness Level (TRL) of 6 for future lunar and Martian missions respectively [11, 27].

A major hurdle to low-risk safe landing of autonomous craft has been a lack of reliable very high-resolution terrain observations, which are necessary to identify all hazards. For this reason, missions prior to 2008 have essentially landed "blind," i.e. without full knowledge of lander-scale hazards from prior orbital observations and without the benefit of on-board hazard avoidance systems. More recently, advances in sub-meter imaging from high-resolution cameras on board reconnaissance orbiters have resulted in surface images with unprecedented detail. Notable are the HiRISE camera on board the Mars Reconnaissance Orbiter (MRO) with a nominal ground sampling distance (GSD) of 0.3m [28], and the Narrow Angle Cameras (NAC) on board the Lunar Reconnaissance Orbiter (LRO) with a nominal GSD of 0.5m [4]. Automated analysis of the images acquired by these electro-optical (EO) sensors have significantly improved landing risk assessment for recent Mars landed missions, namely Phoenix (PHX) [1, 10] and the Mars Science Laboratory (MSL) [11]. Similar analysis is possible for future lunar landed missions [16]. Nevertheless, even with superior orbital observations, PHX and MSL still required statistical approximations of unseen or undetected hazards from orbital analysis. Thus, added capability for on-board hazard detection will be advantageous and minimize risk even for missions whose landing sites have been mapped in high-resolution from orbit.

978-1-4673-1813-6/13/$31.00 ©2013 IEEE

Many scientifically interesting sites tend to be located in terrain that is very hazardous for landing. This compels landings to reasonably benign sites far away from target sites, which may take months to reach. Undoubtedly, adding HDA capability to landed or touch-and-go missions can enable closer access to rougher terrain with high scientific value and can decrease the need of statistically extrapolate the presence of smaller hazards unseen from orbit.

Moreover, high-resolution EO sensors in orbit may not have sufficient sensitivity for sub-meter observations of certain poorly illuminated sites of interest, such as the Moon's poles. ALHAT aims to provide access, regardless of illumination, to these sites via an HDA system that relies on lidar sensors. Also, sub-meter level observations may not be feasible for some distant celestial bodies, such as Europa or Enceladus, due to the cost in funds and time to deploy orbiters there. Missions to such bodies are also likely to benefit from on-board HDA. Furthermore, in a manner similar to the Viking missions to Mars in the 1970s, a modest orbiter can acquire low-resolution imagery for rough landing site selection, and then a lander with HDA can be deployed to choose a safe landing site in real-time.

For Earth applications, on the other hand, the available resources from orbit and ground instrumentation are substantial. Nevertheless, landings on unprepared terrain, whose topography may have undergone significant and sudden changes due to a natural disaster or conflict, requires new analysis of the local terrain prior to attempting a landing. Often there is no sufficient time available for such analysis since these events typically require an immediate response for rescue or assistance operations, day and night. Civilian and military organizations that carry out such operations could, for example, quickly deploy a swarm of small unmanned Short take off and Landing (STOL) craft equipped with HDA and thus having the ability to autonomously choose a safe landing site as close as possible to the target area. An example of efforts in this direction, specifically aimed at unmanned autonomous helicopter landing capabilities, is underway sponsored by the Office of Naval Research (ONR) under the Autonomous Aerial Cargo/Utility Systems (AACUS) program [5].

In this paper, we present an HDA algorithm that analyzes, in real-time, a high-resolution digital terrain model reconstructed during the terminal phase of spacecraft descent. The algorithm first examines the reconstructed terrain to determine the landing safety probabilities of locations on the surface, and then selects and ranks landing sites that exceed tolerable landing risk. The algorithm makes explicit use of the lander's geometry and mechanical tolerances in order to model the touchdown event with increased fidelity. For hazard assessment, the algorithm uses a model of the sensor noise to deal probabilistically with the uncertainty in the terrain measurements. The algorithm also incorporates the navigation uncertainty to adjust the landing safety probabilities accordingly. Lastly, the algorithm selects and report ranked safe landing sites to the spacecraft's Automatic Flight Manager (AFM) for final controlled descent.

We first review previous work related to landing autonomy, then give a brief description of the ALHAT system in Section 2 and Section 3 respectively. Section 4 describes the general HDA algorithm design. Section 5 gives the mathematical foundation for our probabilistic method of modeling landing safety. Section 6 illustrates performance evaluation results. Section 7 describes future work in the implementation and testing of the real-time HDA system. Section 8 gives the concluding remarks.

2. PREVIOUS WORK

Autonomous landings using real-time hazard detection and avoidance systems have not been attempted yet by any planetary mission. However, a number of HDA algorithms have been proposed in the 1990's and early 2000's. These algorithms considered using descent images from electro-optical (EO) passive sensors in monocular and stereo configurations, as well as 3-D point clouds form active scanning and flash Lidar sensors [2, 5, 16, 22, 30, 34, 35]. More recently, technologies for automated detection of rocks and craters from EO imagery were selected by NASA's New Millennium program but a funding shortfall precluded further development [18, 31, 32]. A recent study recommended active sensing as the best choice for a lunar mission due to its ability to work under any sun illumination [3]. In 2007, the ALHAT project began an end-to-end technology development, including HDA, using a Lidar [11]. The first generation HDA algorithm was produced by ALHAT and has gone through substantial sandbox testing, extensive Monte Carlo simulations, and several field tests [20, 25, 28, 29]. Results from these efforts indicated overall good performance, but suggested that further improvements were necessary. In this paper, we address these improvements via the use of a probabilistic framework.

An earlier approach to a probabilistic HDA considered EO single and stereo images to detect rock and slope hazards during the final descent stage [19, 32]. It employed a mathematical model that incorporated the measurement uncertainty and predicted the probability of a safe landing by estimating the probability of finding a region where no hazards are detected times the probability that the selected region has no missed detections. Since the landing scenario was Mars, the availability of a statistical model for the size-frequency distribution of rocks was used to estimate the probability of impacting a hazardous rock as a function of the lander's size [12, 13]. This work suggested that, even under very conservative assumptions about the performance of a vision system, having on-board HDA will reduce the probability of a failed landing by at least a factor of four compared to a blind landing for any rock abundance. Thus, for the level of safety desired by the Mars Science Laboratory (MSL) rover "Curiosity," the vision HDA system would allow access to nearly triple the fraction of the planet as a blind landing would. This would represent a

major improvement in access to sites of scientific value for a small increase in sensor payload. Analogous benefits should accrue to missions to the Moon and to other bodies in the solar system.

3. THE ALHAT SYSTEM

ALHAT's aim is to develop and mature an end-to-end TRL 6 system that uses a Lidar to sense the terrain in real-time and performs safe pinpoint landing anywhere on the Moon [11]. The main processing steps in the ALHAT system are illustrated in Figure 1. Following a terrain relative navigation stage (TRN) [24, 25], a hazard detection and avoidance (HDA) system reconstructs the topography and characterizes the safety of the surface in order to determine the safest landing locations. Then, HDA provides these locations to the Automatic Flight Manager (AFM) for divert and final touchdown maneuvers. During this final descend, a hazard relative navigation (HRN) stage guides the lander to the chosen landing site while keeping navigation drift within bounds [10]. Algorithms for all these stages were incorporated into a software package called Terrain Sensing And Reconstruction (TSAR) [23]. In this paper, we describe the HDA algorithm, encompassing hazard detection from noisy sensor data, consideration of navigation uncertainty, and selection of safe site. All the work reported here has been carried out under the ALHAT project.

The ALHAT effort initially focused on capabilities needed for manned lunar missions in NASA's Constellation program. The requirements called for hazard detection and avoidance on board the proposed Altair crewed lunar lander, a 15-meter diameter spacecraft that had a 0.3m roughness tolerances and a 10^o slope tolerance [4, 37]. Recently, after the cancellation of Constellation, ALHAT is considering the smaller Morpheus lunar lander [17]. The ALHAT effort includes participation by NASA's Johnson Space Center, NASA's Langley Research Center, the C.S. Draper Laboratory, The Johns Hopkins University and the Jet Propulsion Laboratory. Support for field-testing has been provided by NASA's Dryden Center, and more recently, by NASA's Kennedy Space Center.

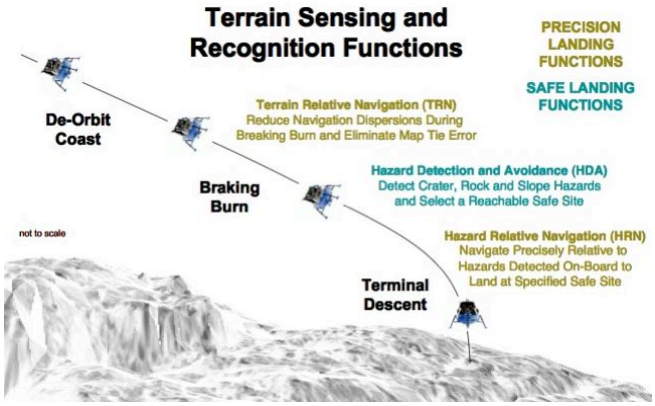

Figure 1 – The ALHAT system

4. HAZARD DETECTION ALGORITHM

4.1 Prior Version

A previous method developed in ALHAT fitted a plane the size of the lander to the terrain in order to determine slope and roughness [20]. However, plane fitting smoothens out roughness features and underestimates the slope of the lander. The proposed algorithm models the interaction between the terrain and the lander, as a function of its geometry and mechanical characteristics, thus, taking full advantage of the high-resolution sensor. By doing so, our HDA algorithm is able to detect the large tilt when a vehicle's pad steps inside a crater or on top of a rock. Also, we can detect roughness due to crater rims and rocks present on the crater rims, where the previous version had trouble.

Also, the previous algorithm used thresholds on the noisy roughness measurements to determine hazards. Due to the elevation errors, the detection thresholds were set lower than the hazard tolerance in order to achieve high hazard detection rates. However, this process introduced large number of false alarms and reduced the available safe area. To deal with this problem, we abandoned the use of specific hazard detection thresholds, and instead implement a probabilistic method to deal more rigorously with the measurement uncertainty associated with the sensor.

4.2 Geometry-based Design

The need to model the touchdown event with greater fidelity was recognized during the Apollo program. Figure 2 shows the Apollo 15 lander straddling the rim of a small hazardous lunar crater. The landing resulted in a vehicle tilt of about 11^o, only 1^o away from the maximum allowable limit, and damaged the engine bell. This incident illustrates the need to estimate the slope of the lander's deck more precisely.

Recently, high resolution LRO imagery allowed detection and mapping of lunar craters as small as 2.5m and boulders larger than 1.5m [7]. Also, Digital Elevation Models (DEMs) with 1.5m postings have been constructed from NAC image stereo pairs [36]. These products provide far improved knowledge of lunar topography and are good tools to use in mission planning. However, they are not sufficient for the touchdown modeling necessitated above. The 0.1m resolution DEM, generated onboard of the ALHAT system, enables high-fidelity modeling and measurements. First, it enables taking into account the lander's geometry explicitly when modeling its contact with the surface at touchdown. Second, it enables calculating effective slope of the lander, i.e. the slope of the plane given by the landing pads. And third, it enables assessing surface roughness of small hazardous features. By definition surface features become landing hazards as a function of the lander's geometry and mechanical slope and roughness tolerances.

Figure 2 – Apollo 15 lander near a crater

The HDA algorithm described here uses the parameters of the Altair lunar lander but without loss of generality. The principles of the algorithm can be used for any lander geometry. It is straightforward to adapt the algorithm for a legged lander of any size with any number of legs, as long as the fidelity of the DEM allows for the modeling of surface contact. The algorithm can be also used for a flat-bottom lander or a non-circular lander.

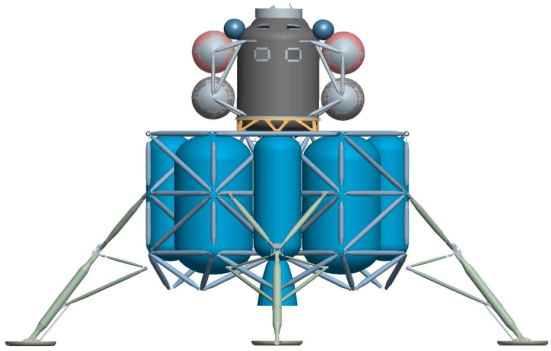

Figure 3 – Altair lander

The Altair lander, illustrated in Figure 3, has a circular footprint of approximately 15m in diameter and is supported on four symmetrically placed legs with 1.5m circular landing pads. We determined the elevation of the pad after surface contact by fitting a robust plane to the DEM area under the pad. This procedure has the effect of averaging out the DEM noise and of simulating any terrain compliance upon contact. The larger the footpad, the less sensitive to noise is the calculated touchdown position. The footpad map, derived from the DEM in Figure 4(a) by local plane fits to each pixel, is illustrated in Figure 4(b). One can see a mound the size of the pad around each rock that the pad steps on. For actual mission design, a higher fidelity model of the pad-terrain contact can be developed, provided there is enough on-board processing time and sufficient DEM resolution.

Since the exact orientation of the lander at touchdown is unknown, the algorithm considers all possible rotational configurations of the lander around its center, also called the aim point. However, orientation at touchdown can be predicted and limited with knowledge of the lander's attitude during descent.

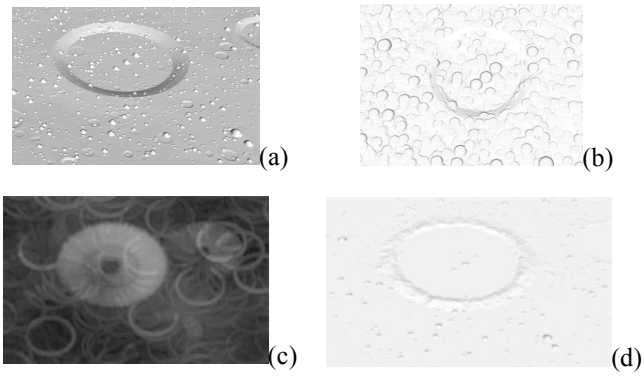

Figure 4 – Portion of a synthetic lunar DEM (a) and the corresponding footpad map (b), slope map (c), and measured roughness map (d).

For a particular rotational configuration of the lander at touchdown, three of the four pads will contact the surface and the remaining pad may be above, or also on, the surface. We assume that the pads are fairly rigid and do not dig deep into the surface. The three pads that contact the surface form a plane whose slope is the effective slope of the lander. This plane is illustrated in blue in Figure 5, with the pad on furthest left corner being in the air. The largest slope from all the orientations is recorded and assigned to the aim point. The results of such computations, performed at each DEM pixel, are shown in Figure 4(c). We can see rings of slope around rocks and crater rims that the lander steps on or in. For the large crater, the lander fits entirely inside and no slope is produced at the center.

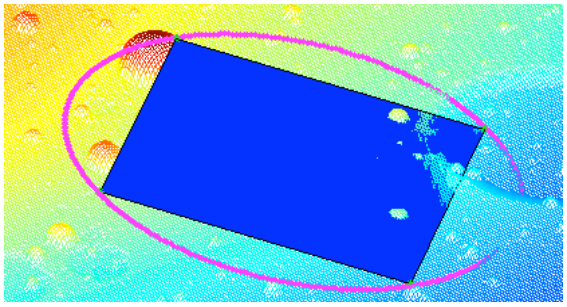

Figure 5 – Lander slope and roughness

The algorithm also determines the presence of roughness under the lander by examining all the terrain under the lander. The largest perpendicular distance from the lander's effective plane to the terrain above it is the value of roughness. The largest roughness is recorded over all orientations. Note that the largest slope and roughness are assessed independently of each other, i.e. they do not have to come from the same lander plane. Figure 4(d) shows the roughness due to rocks and the crater rim. In Figure 5 we can see roughness present on the right hand side.

Finally, the slope and roughness at each DEM pixel is compared against the lander's tolerances to assess a safety for landing.

4.3 Handling sensor noise

The procedure described above is optimal for terrain surfaces reconstructed with very high precision. Yet, surfaces reconstructed from Lidar sensing have inherent sensor noise that makes it difficult to employ such purely deterministic procedures. In our case, the position and slope of the effective lander planes are not affected significantly by noise because it gets averaged out under the large lander pads. For much smaller landers, however, the sensitivity to terrain noise increases and the error in effective slope should be modeled. In any case, the roughness measurements made at individual pixels under the lander, being distances to a noisy terrain, inherit the elevation noise introduced by the sensor.

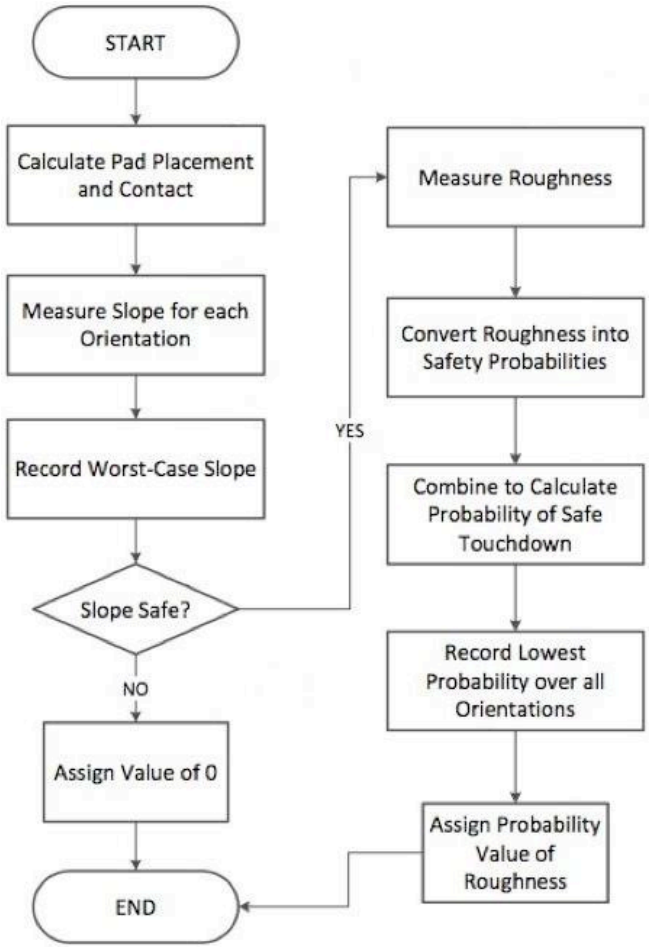

Figure 6 – Assessing safety for an aim point

To better characterize the hazards in the noisy surface, we use a probabilistic model of roughness. Our algorithm computes probability of safety for each pixel under the lander. This probability of safety is the probability that the noisy roughness measurement is above the roughness tolerance given the noise model of the sensor. All output probabilities of safety under the lander are combined together to derive the final probability of safety for each specific landing aim point and orientation. Again, the worst-case probability value is taken across all orientations and is assigned to the aim point. The procedure discussed above is illustrated in Figure 6.

4.4 Handling navigation uncertainty

To incorporate the navigation uncertainty, the algorithm convolves the map comprising of probability values calculated above for all DEM pixels, with a 2-D Gaussian kernel. Thus, the probability of safe touchdown at a particular point inside a landing site is weighted by the probability of going to that point due to navigation error. This procedure produces the final estimate of success at a landing site, or safety score. This safety score correctly captures the effects of sensor noise and navigation uncertainty. For depiction of the entire process see Figure 7.

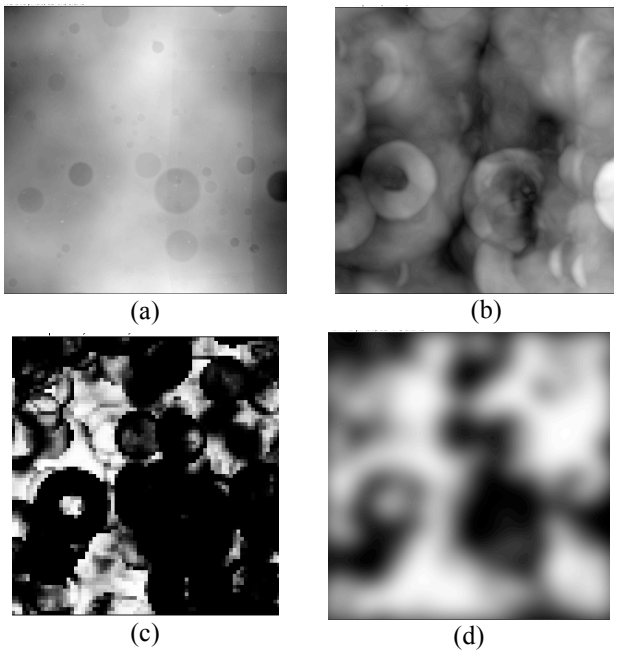

Figure 7– A synthetic Lidar lunar DEM (a) and the corresponding slope map (b), probability of roughness map (c), and probability of safety map (d).

4.5 Safe site selection

Regional maxima are computed in the map resulting from the convolution in 4.4 by applying a morphological domes method. The locations of these maxima represent the landing sites the algorithm reports, ranked in order of decreasing safety. Since the map of scores in 4.4 is dense, this procedure aims to pick sites in different regions of the DEM. The coordinates of the best sites, along with their safety probability, are reported as candidates for landing to the AFM, which in turn, selects the final landing site. Provided fuel and other considerations are met, this final site should be the least hazardous site on the surface.

4.6 Optimizations

Algorithmic optimizations made for real-time operations include discretizing the roughness to centimeter levels and then using pre-computed table of corresponding probability

values, conservatively grouping neighboring roughness values in 5x5 or 10x10 pixel windows, excluding some infeasible orientations, sparsely computing candidate landing sites at 1m intervals, and using separable Gaussian kernels in convolution with the navigation uncertainty. Furthermore, a real-time parallelized version of the HDA algorithm was implemented in C to run on the Tilera Tile64 multi-core processor [38, 39].

5. PROBABILISTIC MODEL

We employ a probabilistic approach to handle the uncertainties in the terrain map elevations and in the navigation system position estimates. Our aim is to achieve a more accurate representation of the safety of the landing site that is not as overly conservative as thresholding. This will allow us to distinguish the safer sites when the noise level is close to the tolerance.

Suppose the lander lands at a particular surface location in a particular orientation. We want to calculate the probability of safe landing given our noisy Lidar DEM. Assume we have already checked the slope of the lander which is largely unaffected by the noise. Then, our hypothesis for safe landing is that, in the true surface, no roughness under the lander exceeds its mechanical tolerance. Let *H* be the set of all possible surfaces satisfying this hypothesis, i.e. all safe surfaces. Then, the probability of safe landing is the probability that the true surface is actually one of the safe surfaces in *H*. For a given possible true surface *S* in *H*, we calculate the probability that *S* could actually be the true surface given the observed Lidar DEM *D* using Bayes' theorem. Then, we sum up the resulting probabilities over the mutually exclusive surfaces of *H*.

$$P(safe \mid D) = \sum_{S \in H} \frac{P(D \mid S)P(S)}{P(D)} = k \sum_{S \in H} P(D \mid S) \quad (1)$$

We assume that all surfaces in *H* are equally likely to be the true surface, i.e. to occur in the natural planetary terrain. Under our assumption, the term *P(S)* is identical for any *S*. The term *P(D)* is also constant. So we let their quotient equal *k*. In reality, however, surface pixels are correlated and some terrains with random jumps in elevation do not occur naturally.

Now consider the discrete pixels making up the true surface. We assume that its pixels have no correlation and thus independently contribute to safety. We further assume that the sensor noise is independent and identically distributed (i.i.d.) at each pixel. So the pixels in the Lidar DEM are also independent. These assumptions let us apply Naïve Bayes rule and computing the probability of safety independently at every pixel. Let *U* be the set of all independent pixels *(x,y)* that are under the lander. A terrain is safe to land if and only if all points *(x,y)* are safe. Thus, the probability of safe landing is the product (logical and) of the probabilities of the individual pixels being safe. In turn, the probability of a given pixel *(x,y)* being safe is the sum (logical or) of the probabilities of it belonging to each one of the safe surfaces.

$$P(safe) = k \prod_{(x,y) \in U} \sum_{S(x,y) \in H} P(D(x,y) \mid S(x,y)) \quad (2)$$

These assumptions are conservative and represent a worse scenario then truth but they make our math practical to implement. Instead of taking the product of the individual probabilities, one could employ a more elaborate method with autocorrelation to determine the overall probability of safety.

For a pixel in *U* to be safe, the perpendicular distance from the lander's actual pad plane to that pixel in the true surface (true roughness) must not exceed the roughness tolerance. However, we only have the noisy Lidar DEM to make this judgment. As we discussed earlier, the calculation of the lander's plane using the Lid r DEM is not sensitive to the Lidar noise present and can be accepted as being the actual value. On the other hand, the distance computed using the Lidar DEM (measured roughness) has error due to elevation noise. The error in the measured roughness at a pixel is proportional to the error in the observed elevation at that pixel.

$$[D(x,y) - S(x,y)] \propto [RD(x,y) - RS(x,y)] \quad (3)$$

To prove that consider, a plane of the lander with equation:

$$ax + by + cz + d = 0 \quad (4)$$

The roughness *r* is computed by calculation the distance to the plane from a point *(x,y,z)* on the surface:

$$r = \frac{ax + by + cz + d}{\sqrt{a^2 + b^2 + c^2}} \quad (5)$$

The error in measured as compared to true roughness is:

$$\begin{aligned} &[RD(x,y) - RS(x,y)] = \\ &\frac{ax + by + cD(x,y) + d}{\sqrt{a^2 + b^2 + c^2}} - \frac{ax + by + cS(x,y) + d}{\sqrt{a^2 + b^2 + c^2}} \\ &= \frac{c}{\sqrt{a^2 + b^2 + c^2}} [D(x,y) - S(x,y)] \end{aligned} \quad (6)$$

Furthermore, the probability of the observed elevation *D(x,y)* given a particular true elevation *S(x,y)* is equal to the probability of the measured roughness *RD(x,y)* given the corresponding true roughness *RS(x,y)*. Thus, we can restate the probability of safety in terms of the roughness at each pixel. The safe roughness values corresponding to the safe elevations $S(x,y) \in H$ ranges from negative infinity up to the roughness tolerance *Tol*. So, at each pixel, we integrate over all safe roughness values the probability that, given a

particular true roughness *r*, the measured roughness *RD(x,y)* occurred.

$$P(D(x,y) \mid S(x,y)) = P(RD(x,y) \mid RS(x,y)) \qquad (7)$$

$$P(safe) = k \prod_{(x,y)\in U} \int_{-\infty}^{Tol} P(RD(x,y) \mid r)\, dr \qquad (8)$$

The elevation noise is modeled as a zero-mean Gaussian with standard deviation σ derived from the calibration of the Lidar instrument. As mentioned above, the elevation noise is assumed to be i.i.d. at each DEM pixel. Under this simple Gaussian model, we have ignored the Lidar's pixel-to-pixel noise variability as well as the change in noise due to the varying incidence angle of the sensor's rays with the surface. Instead, we took the worst-case scenario of direct incidence and maximum pixel noise. In practice, it would be hard to determine the exact noise model for each DEM pixel where the true terrain is not known a priori and the DEM is assembled from mosaicked images.

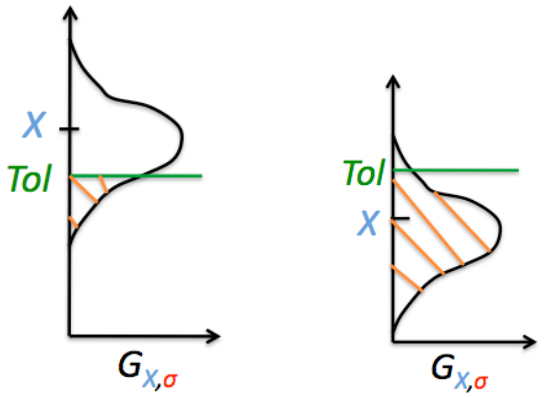

Figure 8 – Probability of roughness

Given our elevation noise model, the error in the observed elevation at a pixel is distributed as a random Gaussian noise. Because of the proportionality, the roughness measured at every pixel is also distributed as a Gaussian, but centered about the true roughness *r* and with standard deviation equal to a constant times σ. Thus, the probability of the measured roughness *rd* given the true roughness *r* is computed by evaluating a Gaussian with mean *r* and sigma σ at *s*. Because of symmetry of the Gaussian, we can switch the mean and evaluation point.

$$\begin{aligned} &[D(x,y) - S(x,y)] \sim \mathrm{N}(0,\sigma^2) \\ &\Rightarrow D(x,y) \sim \mathrm{N}(S(x,y),\sigma^2) \end{aligned} \qquad (9)$$

$$\begin{aligned} &[RD(x,y) - RS(x,y)] \sim \mathrm{N}(0,c^2\sigma^2) \\ &\Rightarrow RD(x,y) \sim \mathrm{N}(S(x,y),c^2\sigma^2) \end{aligned} \qquad (10)$$

$$P(rd \mid r) = \mathrm{N}(r,\sigma^2)\big|_{rd} = \mathrm{N}(rd,\sigma^2)\big|_{r} \qquad (11)$$

This is illustrated in Figure 8, where *X* is the measured roughness, *Tol* is the tolerance, and the black curve is a Gaussian with standard deviation σ equal to that of the noise. The shaded are is the probability of safety. One can observe that even if the measured roughness is slightly above the tolerance, there is still a chance this pixel can be safe due to the noise. On the other hand, a pixel with a measured roughness below the tolerance has a great probability of safety, but a slight chance of hazard.

Substituting the Gaussian expression into the safety equation yields:

$$P(safe) = k \prod_{(x,y)\in U} \int_{-\infty}^{Tol} \mathrm{N}(RD(x,y),c^2\sigma^2)\big|_r\, dr \qquad (12)$$

We can take the constant *c* out:

$$P(safe) = kc^{|U|} \prod_{(x,y)\in U} \int_{-\infty}^{Tol} \mathrm{N}(RD(x,y),\sigma^2)\big|_r\, dr \qquad (13)$$

We observe that the integral expression is simply a Gaussian CDF evaluated at the tolerance:

$$P(safe) = kc^{|U|} \prod_{(x,y)\in U} \Phi\left(\frac{Tol - RD(x,y)}{\sigma}\right) \qquad (14)$$

Note that the proportionality constant *c* is different for each orientation, since the equation of the lander plane changes. The constant *k* is the same for all orientations and all landing locations.

We compute the probability *P(safe)* for every possible orientation, and then we simply take the smallest probability as the final chance of landing safely at a particular location. This conservatively assumes that the lander will always land in the worst-case orientation. A weighted average of the probabilities can be taken instead if the chance of ending up in particular orientation is equally likely.

So far we assumed any random surface is equally possible. One could use a prior for the possible true surfaces. Such prior can be obtained from lower-resolution orbital data and from statistical exponential rock distribution models.

6. Experimental Results

Several test cases were designed to quantify the ability of the algorithm to detect a safe landing site in the presence of sensor noise. Synthetic terrain examples were constructed to be representative of challenging scenarios. The TSAR

package mimicked a 30° trajectory approach to those terrains during which overlapping flashes were collected, starting at 750m slant range, and a Lidar DEM was generated from them. The sensor flashes were simulated using a high-fidelity flash Lidar model provided by NASA Langley Research Center [3]. A Lidar with 256x256 detector and 1° FOV, mounted on a gimbal platform, was assumed. The Altair vehicle geometry with hazard tolerance of 30cm and slope tolerance of 10° was used for the HDA analysis. Monte Carlo trials were run by varying the random Lidar noise.

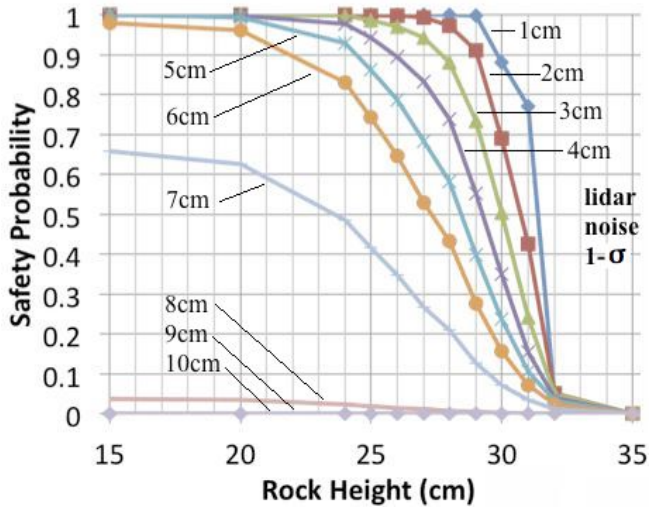

Figure 9 – Effect of Noise on Hazard Assessment

The first test examined the sensitivity of the algorithm's safety probability output to rock height, for various Lidar noise levels. A single rock with varying height was placed in the center of a small DEM the size of the lander. The standard deviation (σ) of the Lidar noise was varied from 1cm to 10cm and 100 trials with different noise realizations were run for each σ value. A curve depicting the mean safety probability versus rock height is plotted for each Lidar noise level in Figure 9. Examining the family of curves from top to bottom, one can see that as the Lidar noise increases, the safety probabilities for rocks decreases. This effect is most apparent for rocks of heights between 20cm and 32cm depicted in the middle vertical section of the plot. Thus, in the presence of Lidar noise, the HDA algorithm is most sensitive to rocks of these heights, which are spread around the vehicle's tolerance. Examining an individual curve from left to right, one can see that the safety probability monotonically drops as the rock height increases. In general, the probability starts dropping for rocks larger than 25cm and becomes virtually zero for rocks larger than 35cm. The sleeper the knee in the curves is around the tolerance, the easier it is to distinguish hazards from non-hazards since the average safety of a small hazardous rock will be much lower than that of a large safe rock. One can see that our HDA algorithm make such distinction well to the Altair lander for Lidar noises up to 5cm 1-σ. Note, however, that all safety probabilities outputted by HDA drop significantly for Lidar noises of 8cm 1-σ or greater. Given the size and tolerances of the Altair vehicle, our HDA algorithm will not work for a Lidar sensor with such large amount of noise. Nonetheless, additional experiments confirmed that our algorithm would function in the presence of 8cm Lidar noise if the vehicle's hazard tolerance were 40cm.

The second test examined the sensitivity of the algorithm's safety probability output to rock count for varying rock heights. This test compared the safety probability between landing sites with a varying number of rocks of certain heights ranging from 15cm to 35cm. For a set Lidar noise of 5cm 1-σ, Figure 10 shows a family of curves, one for each rock count, plotting safety probability versus rock height. As the number of rocks of a certain height increases, the safety probability decreases. There is a point when a large number of non-hazardous rocks appear less safe than a single hazardous rock due to noise. To illustrate this, a black reference line corresponding to a site with one hazardous rock of 30cm in height is plotted. Observe that two 27cm rocks appears less safe than one 30cm rock and so do four 26m rocks. Nevertheless, this confusion of a safe site containing marginally safe rocks (25-30cm) with an unsafe having one minimally hazardous rock is rare to occur. Lunar rock abundance suggest that on Smooth Mare only 2 to 3 rocks greater than 25cm in height will occur per Altair-size area [18]. Fundamentally, our algorithm can distinguish between sites with number of non-hazards and sites with a few hazards. Even further, landers with smaller footprints will encounter fewer rocks. Thus, given the same sensor noise and hazard tolerance, smaller landers will have no problem landing on Smooth Mare using our algorithm.

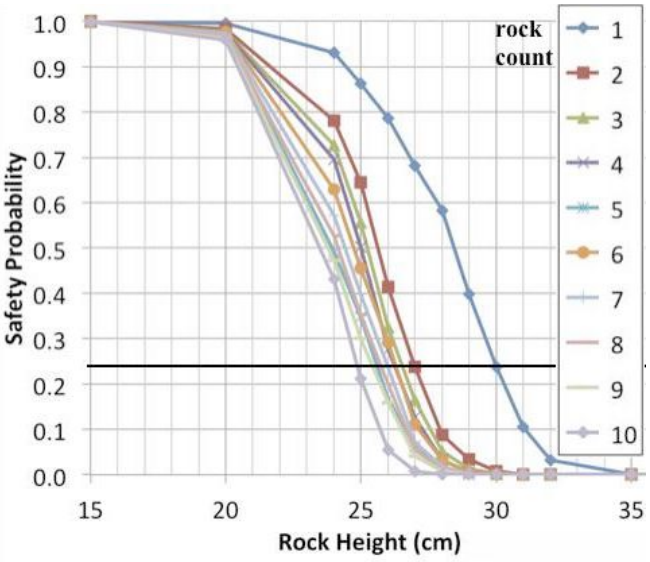

Figure 10 – Effect of Rock Abundance on Hazard Assessment

The third test examined the ability of HDA to rank landing sites in the presence of Lidar noise and navigation uncertainty. For this purpose, we constructed a DEM with four quadrants of known safety separated by hazardous

walls as shown in Figure 11. The quadrants and their rankings were as follows: completely flat and safe quadrant (rank 1); quadrant of rocks with heights less than 25 cm also safe (rank 2); quadrant of rocks with mixed heights containing one safe site (rank 3); quadrant of all hazardous rocks containing no safe site (rank 4). 100 trials were run for each combination of 5, 8, and 10cm 1-σ Lidar noise and 0, 25, 30, 35, 40, 45, 50cm random trajectory jitter. The HDA algorithm reported a ranked list of landing sites within the DEM. Only in 6 of the total 2100 trials the ranking was different than expected. In those cases, the safe quadrants 1 and 2 were swapped, due to excessive noise and jitter. Therefore, we are very confident that the algorithm can rank properly evidently different sites.

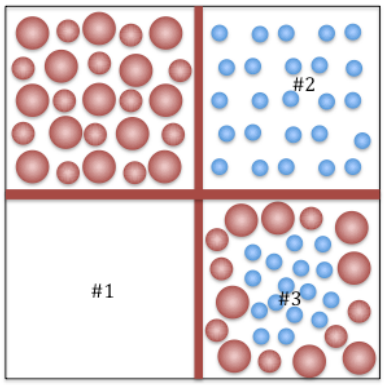

Figure 11 – Cartoon version of the Digital Elevation Map with Four Quadrants of Known Safety Rank

For the fourth test, lunar terrain was simulated using a DEM generation procedure developed at JPL that creates fractal terrains and adds rocks and craters from model distributions [18]. Ten synthetic lunar-like Smooth Mare terrain maps were generated and each was populated with increasing rock abundance of 0, 2, and 4%. 200 trials were run on each of the 30 maps with 5cm 1-σ Lidar noise. A Gaussian navigation uncertainty with 1m 1-σ was assumed. All top 5 HDA selected landing sites for the 0% and 2% rock abundance terrains were at least 99% safe in the reference. For the 4% terrains, on average 4.5 of the top 5 and 2.8 of the top 3 sites were 99% safe. The top 1 landing site was always at least 99% safe. This study shows that our algorithm should always find a safe site in Smooth Mare lunar terrains with 4% or less rock abundance.

The fifth test involved studying the slope estimation of the HDA algorithm. Similarly to test four, we generated 50 maps for each of these two lunar terrain types: Smooth Mare (SM), the most benign, and Hammock Uplands (HU), the most hazardous. We ran 50 trials on each map with 5cm 1-σ Lidar noise. Results showed that the maximum absolute error for slope was 0.045° on SM and 0.14° on HU. The average error was 0.018° on SM and 0.05° on HU. These insignificant errors show superb performance in estimating the effective lander slope in the presence of noise and thus in detecting slope hazards.

7. Future Work

A great deal of engineering work is planned and is currently being carried out by the ALHAT team. The integration of the system hardware and software components is progressing through a number of test campaigns inside the laboratory and in the field. Particularly, a final integration test campaign is scheduled at the Long Distance Range Test Range (LDTR) at NASA's Langley Research Center in the late fall of 2012.

Recently, a one-hectare hazard field illustrated in Figure 12 was constructed at NASA's Kennedy Space Center (KSC) for testing and evaluation of ALHAT's autonomous landing technologies. The field's features were based on a size-frequency distribution of boulders and fresh craters derived from an actual 500 by 500 meter lunar terrain patch with 3.4% rock abundance. A flight on-board a helicopter over this hazard field is planned for the winter of 2012.

Concurrently with ALHAT, the Morpheus team at NASA's Johnson Space Center has been focusing over the last several months on tether testing of the Morpheus prototype lunar lander [17]. This lander is much smaller than Altair, with a 3.4-meter diameter. A joint ALHAT-Morpheus test campaign over the hazard field at KSC is scheduled in the spring of 2013. The goal of this test will be to demonstrate autonomous safe landing of the prototype lander guided by the on-board ALHAT system. Successful completion of this test campaign would bring the technology to TRL level 6, the intended goal of the ALHAT project.

Figure 12 – Hazard Test Field at NASA's Kennedy Space Center

8. CONCLUSION

We have developed and presented a new approach for the design and implementation of a real-time hazard detection and avoidance system for autonomous safe landing. Our approach has four main contributions. First, we incorporate the lander's geometry and relevant mechanical tolerances to allow higher fidelity in the modeling of the touchdown event. High-resolution measurements of the surface topography from a Lidar sensor enable us to estimate well the effective lander slope and roughness. Second, we use a probabilistic formulation to deal explicitly with the sensor noise that entirely eliminates the need to set deterministic thresholds for hazard assessment. Given the lander's tolerances to surface features and a model of the sensor noise, the probability of a hazard under the lander's deck is determined analytically and robustly. Third, we include the expected navigation uncertainty into the estimation of the safe landing probabilities associated with each possible landing sites in the DEM of the terrain. Furth, we detect and rank safe sites to enable the AFM, or a crewmember, to execute a divert maneuver. A patent is pending for the general design of the HDA algorithm.

A significant effort has been devoted to testing this approach within the integrated ALHAT system in order to demonstrate its utility, in real-time, for future landed missions. Preliminary studies on synthetic terrain indicate that fundamentally our algorithm can distinguish roughness hazards in the presence of Lidar noise. Also, our algorithm has shown superb performance in estimating the effective lander slope in the presence of noise and thus in detecting slope hazards. Experiments show that the HDA algorithm can reliably find a safe landing site for the Altair lander on Smooth Mare lunar terrain.

ACKNOWLEDGEMENTS

The work described in this publication was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract from the National Aeronautics and Space Administration. The authors would like to thank the ALHAT Team at JPL and in particular Andrew Johnson and John Carson, the former and current ALHAT Program managers for JPL, for help with algorithm development and analysis of results. Also, we thank the ALHAT teams at the NASA's Johnson Space Center, NASA's Langley Research Center, NASA's Dryden Center, the C.S. Draper Laboratory, and NASA's Kennedy Space Center. This work was performed under the ALHAT Project managed by NASA Johnson Space Center and funded by the NASA Exploration Technology Development Program.

REFERENCES

[1] Arvidson, R. E., et al., "Mars Exploration Program 2007 Phoenix landing site selection and characteristics," *J. Geophysical Research*, Vol. 114, No. E3, June 2008.

[2] Bajracharya, M., "Single Image Based Hazard Detection for a Planetary Lander," *Proc. of Automation Congress*, Vol. 14, pp. 585-590, June 2002.

[3] Brady, T. and J. Schwartz, "ALHAT System Architecture and Operational Concept," *Proc. Aerospace Conf.*, Big Sky, MT, Mar. 2007.

[4] Brady, T. et al., "Hazard Detection Methods for Lunar Landing," *Proc. Aerospace Conf.*, Big Sky, MT, Mar. 2009.

[5] Bulyshev, A. et al., "A Super Resolution Algorithm for Enhancement of Flash Lidar Data," *Proc. of SPIE Computational Imaging IX*, Vol 7873, Feb. 2011.

[6] Cheng, Y., A. Johnson, L. Matthies, A. Wolf, "Efficient Passive-Image Based Hazard Detection for Safe Landing on Mars," *Int'l Symp. Artificial Intelligence, Robotics and Automation in Space* (iSAIRAS), Montreal, Canada, June 2001.

[7] Cheng, Y., D. Clouse, A. Johnson, W. Owen, and A. Vaughan, "Evaluation and Improvement of Passive Optical Terrain Relative Navigation Algorithms for Pin-Point Landing," *Proc. AAS Space Flight Mechanics Meeting* (AAS-SFM), Jan. 2011.

[8] Chin, G. et. al., "Lunar Reconnaissance Orbiter overview: The instrument suite and mission," *Space Science Reviews*, Vol. 129, No.4, pp.391-419, Apr. 2007.

[9] Cummings, M. and A. Collins, "Autonomous Aerial Cargo/Utility System (AACUS), Concept of Operations (CONOPs)," Office of Naval Research, Dec. 2011, Accessed Oct. 10, 2012, <www.onr.navy.mil/~/media/Files/Funding-Announcements/BAA/2012/12-004-CONOPS.ashx>.

[10] Davis, J. et al, "Post2 End-to-End Descent and Landing Simulation for ALHAT Design Analysis Cycle 2," *International Planetary Probe Workshop*, Barcelona, Spain, June 2010,

[11] Epp, C., and T. Smith, "Autonomous Precision Landing and Hazard Detection and Avoidance Technology (ALHAT)," *Proc. of Aerospace Conf.*, Big Sky, MT, Mar. 2007.

[12] Golombek, M., and D. Rapp, "Size Frequency Distribution of Rocks on Mars and Earth Analog Sites; Implications for Future landed Missions," *J. Geophysical Research*, Vol. 102, No. E2, pp. 4117-4129, Feb. 1997.

[13] Golombek, M. et al., "Rock Size-Frequency Distributions on Mars and Implications on Mars Exploration Rover Landing Safety and Operations," *J. Geophysical Research*, Vol. 108, No. E12, pp. 8086-8109, Oct. 2003.

[14] Golombek, M., A. Huertas, et al., "Size-frequency distributions of rocks on the northern plains of Mars with special reference to Phoenix landing surfaces," *J. Geophysical Research*, Vol. 114, No. E3, July 2008.

[15] Golombek, M., A. Huertas and D. Kipp, "Detection and Characterization of Rocks and Rock Size-Frequency Distributions at the Mars Science Laboratory Landing Sites," *MARS: Int. J. Mars Science and Exploration,* Vol. 7, pp. 1-22, Aug. 2012.

[16] Halbrook, T., J. Chapel, J. Witte, "Derivation of Hazard Sensing and Avoidance Maneuver Requirements for Planetary Landers," *AAS Guidance and Control Meeting*, Breckenridge, CO, Jan 2001.

[17] Hart, J. and J. Mitchell, "Morpheus lander testing campaign," *Proc. Aerospace Conf.*, Big Sky, MT, Mar. 2012.

[18] Huertas, A., Y. Cheng, and R. Madison, "Passive Imaging Based Multi-cue Hazard Detection for Spacecraft Safe Landing," *Proc. Aerospace Conf.*, Big Sky, MT, Mar. 2006.

[19] Huertas, A., Y. Cheng and L. Matthies, "Real-time Hazard Detection for Landers," *NASA Science Technology Conf.*, College Park, Maryland, June 2007.

[20] Huertas, A. Johnson, R. Werner, and R. Maddock, "Performance Evaluation of Hazard Detection and Avoidance Algorithms for Safe Lunar Landings." *Proc. of Aerospace Conf.*, Big Sky, MT, Mar. 2010.

[21] Huertas, A. and Y. Cheng, "Automatic Mapping of Lunar Craters and Boulders," *Proc. of Lunar and Planetary Science Conf.*, The Woodlands, TX, Mar. 2011.

[22] Johnson, A., A. Klumpp, J. Collier, and A. Wolf, "Lidar-based Hazard Avoidance for Safe Landing on Mars," *AIAA J. Guidance, Control and Dynamics,* Vol. 25, No. 5, Oct. 2002.

[23] Johnson, A., A. Huertas, R. Werner, J. Montgomery, "Analysis of On-Board Hazard Detection and Avoidance for Safe Lunar Landing," *Proc. of Aerospace Conf.*, Big Sky, MT, Mar. 2008.

[24] Johnson, A., and J. Montgomery, "An Overview of Terrain Relative Navigation Approaches for Precise Lunar Landing," *Proc. of Aerospace Conf.,* Big Sky, MT, Mar. 2008.

[25] Johnson, A., J. Keim, and T. Ivanov, "Analysis of Flash Lidar Field Test Data for Safe Lunar Landing," *Proc. of Aerospace Conf.*, Big Sky, MT, Mar. 2010.

[26] Johnson, A., and T. Ivanov, "Analysis and Testing of a LIDAR-based Approach to Terrain Relative Navigation for Precise Lunar Landing," *Proc. of AIAA GNC Conference*, Portland, OR, Aug. 2011.

[27] Johnson, A., and M. Golombek, "Lander Vision System for Safe and Precise Entry Descent and Landing" *Concepts and Approaches for Mars Exploration*, Houston, Texas. June 2012.

[28] Keim, J., et al., "Field Test Implementation to Evaluate a Flash Lidar as a Primary Sensor for Safe Lunar Landing," *Proc. of Aerospace Conf.*, Big Sky, MT, Mar. 2010.

[29] Keim, J., S. Mobasser, E. Bailey, A. Johnson, and G. Khanoyan, "Test Implementation to Evaluate Technologies for Safe Lunar Landing" *Proc. of Aerospace Conf.*, Big Sky, MT, Mar. 2011.

[30] Langley, C. et al., "Recent Advancements of the Lidar-Based Autonomous Planetary Landing System." *Proc. of Int'l Astronautical Conference*, Hyderabad, India, Sept. 2007.

[31] Matthies, L., A. Huertas, Y. Cheng and A. Johnson, "Landing Hazard Detection with Stereo Vision and Shadow Analysis," *AIAA Infotech at Aerospace Conf.*, Big Sky, MT, Mar. 2007.

[32] Matthies, L., A. Huertas, Y. Cheng, and A. Johnson, "Stereo Vision and Shadow Analysis for Landing Hazard Detection," *Proc. of ICRA*, Pasadena, CA, May 2008.

[33] McEwen, A. S., et al., "Mars Reconnaissance Orbiter's High Resolution Imaging Science Experiment (HiRISE)," *J. Geophysical Research*, Vol. 112, No. E5, May 2007.

[34] Pien, H. "Autonomous Hazard Detection and Avoidance," *C.S. Draper Laboratory Report*, CSDL-C-6303, June 1992.

[35] Pien, H., G. Barton, T. Brand, and D. Suh, "Assessment of Blind Landing Risk on Viking Specific Terrain," *C.S. Draper Laboratory Report*, CSDL-R-2688, July 1995.

[36] Rosiek, M. et al., "USGS Digital Terrain Models and Mosaics for LMMP," *43rd Lunar and Planetary Science Conf.*, The Woodlands, TX, Mar. 2012.

[37] Strahan, A. L., and A. Johnson, "Terrain Hazard Detection and Avoidance during the Descent and Landing Phase of the Altair Mission," *Proc. of AIAA GNC Conference,* Toronto, Canada, Aug. 2010.

[38] Villalpando, C., A. Johnson, J. Oberlin, and S. Goldberg, "Investigation of the Tilera Processor for Real-time Hazard Detection and Avoidance on the AltairLunar Lander," *Proc. IEEE Aerospace Conf.*, Big Sky, MT, Mar. 2010.

[39] Villalpando, C., R. A. Werner, J. M. Carson III, G. Khanoyan, R. A. Stern and N. Trawny, "A Hybrid FPGA/Tilera Compute Element for Autonomous Hazard Detection and Navigation," *Proc. IEEE Aerospace Conf.*, Big Sky, MT, Mar. 2013

BIOGRAPHIES

Tonislav Ivanov *is a Research Technologist II at the Jet Propulsion Laboratory. He received his B.S. and M.Eng. in Electrical and Computer Engineering from Cornell University in 2007. After graduation, he joined the Computer Vision Group at JPL. He works on lidar data processing, field test setup, and development of the hazard detection algorithm for the Autonomous Landing and Hazard Avoidance Project. He also prototyped the hazard detection algorithms for the Lander Vision System of the Mars Technology program. He worked on algorithms for helicopter landing. He also worked on generation of slope and roughness hazard maps for the Lunar Mapping and Modeling Project. His other work at JPL includes pedestrian detection and human face, pose, and activity recognition from stereo image data for robot human awareness.*

Andres Huertas *is a Senior Member of the Technical–Staff in the Computer Vision Group at the Jet Propulsion Laboratory. He earned his B.S. degree in Systems and Computing Engineering from Universidad de los Andes in 1974, and his M.S. and Engineer degrees in Computer Science in 1978, and Electrical Engineering in 1984 from the University of Southern California. Prior to joining the Computer Vision Group at JPL in 2002, he worked on computer vision research since 1978 at the Image Processing Institute and at the Institute for Robotics and Intelligent Systems at the University of Southern California. At JPL, he works on computer vision techniques for autonomous off-road navigation, and on techniques for hazard detection for safe landing of spacecraft on planetary surfaces. He participated in hazard risk assessment for the certification of candidate landing sites for the Phoenix and Mars Science Laboratory Missions. He has been a member of the ALHAT team for several years and participated in the Lunar Mapping and Modeling Project to automatically detect and help characterize the size-frequency distributions of craters and boulders from NAC orbital imagery of the Lunar surface. With co-author Toni Ivanov, they designed the Hazard Field at KSC for upcoming field-testing of ALHAT technology.*